

TITLE OF THE INVENTION**STORAGE MANAGEMENT SYSTEM****BACKGROUND OF THE INVENTION****5 Field of the Invention**

The present invention relates to a storage management system and, in particular, to a volume allocating method in a technology for storage operation management on a Storage Area Network (SAN).

10 Description of the related Art

The SAN effectively uses a large amount of data and, therefore, it is extremely widespread in recent years. The introduction of the SAN increases the flexibility of an information system in a company. However, a plurality of 15 storage devices are connected via a fiber channel switch and this complicates the structure of the storage device. Further, the equipment setting of a system manager becomes complicated.

The central control of the storage devices causes the 20 increase in data capacity of a single system manager. Further, this increases the types of service servers and the types of service application. In addition, the amount of data increases day by day and, therefore, the storage device is enhanced and the storage allocation to the 25 service server is frequently executed.

In order to actually store data in the storage device, a logical area, a so-called volume needs to be allocated to the storage device. The allocation of volume needs not only the designation of storage device but also the 5 designation of disk group in the storage device, namely, group of disks forming RAID structured by collecting a plurality of disk devices.

One well-known technology for central control of the complicated storage devices is disclosed in U.S. Patent 10 Publication No. 6,430,611 (Patent Reference 1). According to the technology disclosed in U.S. Patent Publication No. 6,430,611, a computer for managing a data area in the storage device is connected to the storage devices connected to the SAN, and agents for collecting information 15 on the storage devices operated therein are arranged. A single management server for communication with the agents is provided and, thus, it can manage the entire storage devices in a network (refer to the 56th line in the first paragraph to the 36th line in the second paragraph in U.S. 20 Patent Publication No. 6,430,611).

A plurality of volumes can be arranged in the single disk group and, therefore, the simultaneous access to a plurality of volumes in the single disk group reduces the individual volume performance. One technology for reducing 25 the load of volume in the single disk group is disclosed in

U.S. Patent Publication No. 6,535,954 (Patent Reference 2)
(corresponding to Japanese Unexamined Patent Application
Publication No. 11-296313). According to the technology
disclosed in U.S. Patent Publication No. 6,535,954, the
5 load is concentrated to the disk group and, when one disk
control device has excessive load prior to the load of the
disk group, the management of the disk group or volume
shifts to another disk control device, thereby exhibiting
the maximum performance of the disk group or volume (refer
10 to the 17th to 27th lines in the third paragraph in U.S.
Patent Publication No. 6,535,954).

Further, Japanese Unexamined Patent Application
Publication No. 2001-184175 (Patent Reference 3) discloses
a technology for storing the history of access information
15 to the storage device and for forming a candidate of new
logical-disk-structure based on the analysis result of the
history information so as to restructure the logical disk
structure in which the access performance is improved and
the using efficiency of the disk area is increased.

20 The complexity of the structure of storage device
complicates a volume management method. Therefore, the
solution of the above problem is demanded.

However, U.S. Patent Publication No. 6,430,611
discloses that the single management server can monitor the
25 complicated storage devices. However, a volume management

method is not disclosed.

Further, according to U.S. Patent Publication No. 6,535,954, the maximum performance of the disk group is exhibited. However, there is a problem that a plurality of 5 volumes are arranged in the disk group and the requested performance is not exhibited when the disk group has the load of the maximum performance or more.

When the number of accesses over the theoretical maximum performance of the disk group is requested to a 10 plurality of volumes formed in the disk group, the disk group cannot satisfy the requested performance. Thus, preferably, upon allocating the volume, the disk group can be selected so as to prevent the excessive performance of the disk group by allocating the volume.

15 Further, the allocation of volume for requesting the higher performance needs the efficient allocation of volume so as to remain the disk group for non-allocation as much as possible.

The technology disclosed in Japanese Unexamined Patent 20 Application Publication No. 2001-184175 pays attention to the increasing ratio of the used disk capacity. However, the candidate of the new logical-disk structure is not determined by paying attention to the performance and the reliability. Further, Japanese Unexamined Patent 25 Application Publication No. 2001-184175 does not describe a

method for systematically allocating the volume before operation while anticipating the deterioration in performance.

5

SUMMARY OF THE INVENTION

It is one object of the present invention to prevent the deterioration in performance which is caused by the excessive access performance of the allocated disk group when the volume is allocated and is actually operated.

10 It is another object of the present invention to efficiently allocate the volume based on operation history of the volume so as to prevent the disk group from having its performance which is over the maximum one.

According to the present invention, a storage management server comprises: a database for operation history which stores, as history, information including a performance value of a disk group obtained upon operating the storage device; a database for a volume performance value which stores information on specification values including performance, reliability, and a capacity of the storage device obtained from the storage device; a policy database which stores information on policies including the performance corresponding to a plurality of set policies; first processing means which calculate a forecasted performance value from the information on the performance

value of the disk group stored in the database for
operation history; second processing means which obtain a
performance margin, based on a theoretical performance
value of the volume and the forecasted performance value
5 obtained by the first processing means; and volume
determination processing means which determine an
allocation candidate of the volume in accordance with the
calculation result of the second processing means.

The first processing means calculate the forecasted
10 performance value per unit time based on information on the
performance value obtained from the database for operation
history, and the database for a volume performance value
stores information on the forecasted performance value per
unit time obtained by the first processing means,
15 corresponding to the disk group.

The second processing means perform processing for
obtaining the difference between the performance margin per
unit time and the designated performance value stored in
the policy database, and the volume determination
20 processing means determine, as the allocation candidate,
the volume which is obtained by the second processing means
and has a positive difference.

According to the present invention, the storage
management server further comprises: means which transmit
25 information indicating the volume candidate determined by

the volume determination processing means so as to display the information on a client connected to the storage management server; and means which receive the information on the volume allocation selected by the client in accordance with the displayed information.

Further, according to the present invention, in a storage management system, a client comprises: means which designate and input the condition for allocating the volume; display means that display the information indicating the volume candidate determined by the volume determination processing means; and means which transmit, to the storage management server, the information on the volume allocation selected from the volume information of the allocation candidate displayed on display means.

Furthermore, the present invention can be applied to a program. That is, there is provided a program for selecting and generating a volume candidate functioning on a storage management server having a database on operation history for storing, as history, information including a performance value of a disk group obtained by operating a storage device connected via a network, a database for a volume performance value for storing information on specification values including performance, reliability, and a capacity of the storage device, obtained from the storage device, and a policy database for storing

information on a policy including the performance corresponding to a plurality of set policies. The program for generating the volume candidate comprises: a first processing step of calculating a forecasted performance value from the information on the performance value of the disk group stored in the database on the operation history; a second processing step of obtaining a performance margin based on a theoretical performance value of the volume and the forecasted performance value obtained by the first processing step; a volume determination processing step of determining a candidate for allocating the volume in accordance with the calculation result of the second processing step; and a step of generating information for displaying the volume candidate from information based on the volume determination processing step, so as to display the volume candidate on a client connected to the storage management server.

In addition, the present invention can be applied to a volume allocating method in a storage management system. According to the present invention, there is provided a volume allocating method in a storage management system for managing the operation of a storage device connected via a network by use of a storage management server. The volume allocating method comprises receiving, via the network, a condition for allocating a volume designated by a client;

obtaining information on operation history of the volume from a memory device for storing, as history, information including a performance value of the disk group obtained upon actually operating the storage device; obtaining

5 information on specification values including the performance value of the storage device; assuring a performance margin and determining a candidate of an allocable volume in accordance with the received condition for allocating the volume based on the information on the

10 operation history of the volume and the information on the storage device; transmitting information on the volume of the allocated candidate to the client; receiving information on volume allocation selected and transmitted from the information on the volume of the allocated

15 candidate in the client; and allocating the volume to the storage device in accordance with the information on the volume allocation.

The volume allocating method further comprises:

previously storing, in the memory device, a plurality of

20 policies one of which is selected by designating the condition for allocating the volume in the client, including at least the performance value and an operating time zone; and previously storing, in the memory device, information on a forecasted performance value per unit time

25 which is calculated from a capacity, a theoretical

performance value, and information on the operation history of the volume of the disk group as an allocation target.

The step of determining the volume candidate comprises: obtaining the performance margin based on the theoretical performance value and the forecasted performance value per unit time of the volume included in the disk group; calculating and subtracting the performance value designated by the policy from the obtained performance margin; and a step of determining, as the allocation candidate, the volume of the disk group when the obtained value is positive as the calculation result.

In the volume allocating method according to the preferred embodiment, the new volume is allocated to a disk group and the performance necessary for the disk group is obtained based on information on the operation history of the disk groups and on the operating schedule of the newly-allocated volume. Consequently, it is possible to determine whether or not the disk group satisfies the performance necessary for the allocated volume upon allocating the volume.

Further, in the volume allocating method according to the present invention, the volume as the allocation target is selected based on information on performance necessary for the obtained performance and information on the disk group executable performance. Consequently, upon

allocating the volume which satisfies the newly-requested performance, the disk group for allocation is selected while the disk group having the higher performance remains.

The operating state of volume is not constant but
5 changes depending on services to which the volume is allocated. As a result of using the abovementioned method, a plurality of volumes having the different operating states can be allocated to the single disk group so as to prevent the synchronicity of their using peaks, and the
10 performance necessary for the volume can be satisfied if the disk group has the performance which is lower than the required performance obtained from the peak performance for the services.

15 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram showing the structure of a storage management system used for the SAN and the outline of a processing sequence thereof according to an embodiment of the present invention;

20 Fig. 2 is a policy table of a policy database 35 in the storage management system according to the embodiment;

Fig. 3 is a table of a database 36 for a volume performance value in the storage management system according to the embodiment;

25 Fig. 4 is a table of a database 37 for operation

history in the storage management system according to the embodiment;

Fig. 5 is a flowchart for processing for obtaining a performance value of a disk group according to the 5 embodiment;

Fig. 6 is a flowchart for processing for determining the volume of an allocated candidate according to the embodiment;

Fig. 7 is a diagram showing an example of a display 10 screen for designating a condition for allocating the volume according to the embodiment;

Fig. 8 is a diagram showing the presentation of the volume candidate and a selected screen according to the embodiment; and

15 Fig. 9 is a diagram showing an example of calculating a performance margin according to the embodiment.

DESCRIPTION OF PREFERRED EMBODIMENTS

A detailed description is given of an embodiment of 20 the present invention with reference to the drawings.

Fig. 1 is a diagram showing the structure of a storage management system used for the SAN and the outline for processing sequence thereof.

The storage management system comprises: a storage 25 management client 2; a storage management server 3; a

service host 5; and a storage device 7. These components in the storage management system are connected by a LAN (Local Area Network) 4, and information is referred to and management information is set between the devices via the
5 LAN 4. The service host 5 is connected to the storage device 7 via an SAN 6, thereby receiving and transmitting data used for the service via the SAN 6.

The storage management client 2, such as a personal computer, comprises: a display device for displaying
10 information; an input device, e.g., a keyboard or a mouse; a memory device, e.g., a hard disk device (HDD); a processor; and an internal memory. The memory device prepares and stores a program 20 for presenting a volume candidate. The program 20 for presenting the volume
15 candidate is loaded to the internal memory and is executed by the processor upon allocating the volume candidate. An operating manager 10 inputs data on the allocation of volume candidate to the input device. A display screen of the display device displays information on the volume
20 candidate received from the storage management server 3, and presents the volume candidate to the operating manager
10.

A description is given of functions of the program 20 for presenting the volume candidate on the storage
25 management client 2. The program 20 for presenting the

volume candidate first receives an input of a condition for allocating the volume from the operating manager 10 on processing 21 for designating the condition for allocating the volume, and transmits the designated condition to a 5 program 31 for generating the volume candidate in the storage management server 3.

The program 31 for generating the volume candidate generates the volume candidate and transmits the volume candidate to the program 20 for presenting the volume 10 candidate. In processing 22 for presenting and selecting the volume candidate, the execution of the program 31 for generating the volume candidate is waited and the generated volume candidate is received. The received volume candidate is displayed on the display screen and the volume 15 candidate is presented to the operating manager 10. The operating manager 10 selects a proper volume from the displayed candidates. In processing 23 for allocating the volume, a request for allocating the volume selected and inputted by the operating manager 10 is transmitted to the 20 storage management server 3.

The storage management server 3 has the storage device. Further, the storage management server 3 prepares and stores the program 31 for generating the volume candidate, a program 32 for collecting the operation history, a 25 program 33 for collecting specification values, and the

program 34 for allocating the volume. These programs are executed by a processing device (not shown). Further, the memory device comprises a policy database (DB) 35, a database (DB) 36 for a volume performance value, and a

5 database (DB) 37 for the operation history. Memory formats of the databases 35 to 37 are described in detail later.

Further, the characteristic functions and operations of the program 31 for generating the volume candidate will be described in detail later.

10 The program 31 for generating the volume candidate obtains information from the databases 35 to 37 based on the condition for allocating the volume which is received from the storage management client 2, generates the volume candidate, and sends the generated volume candidate to the
15 storage management client 2. The program 32 for collecting the operation history collects the operation history of the storage device 7, and stores the collected operation history to the database 37 for the operation history. The program 33 for collecting the specification values collects
20 the specification values of the storage device 7, and stores the collected specification values to the database 36 for the volume performance value. The program 34 for allocating the volume receives the request for allocating the volume transmitted from the storage management client 2,
25 and allocates the volumes to the storage device 7.

A description is given of the functions of the program
31 for generating the volume candidate on the storage
management server 3.

In processing 311 for obtaining an allocable disk
5 group, the program 31 for generating the volume candidate
performs processing for allocating the candidate of the
allocable disk group based on the condition for allocating
the volume of the operating manager 10 received from the
program 20 for presenting the volume candidate and on
10 specification values 71 for obtaining from the storage
devices. In processing 312 for forecasting the performance
value of the disk group, the program 31 for generating the
volume candidate forecasts the performance of the disk-
group candidate based on the performance value which is
15 actually obtained from the database 37 for the operation
history, and stores the forecasted performance value into
the database 36 for the volume performance value. In
processing 315 for determining the allocated candidate
volume, the program 31 for generating the volume candidate
20 determines the disk group matching the condition for
allocating the volume designated by the operating manager
10 from the forecasted disk-group candidates. If a
plurality of the determined disk groups exist, in
processing 318 for generating information indicating the
25 volume candidate, the program 31 for generating the volume

candidate arranges the disk groups in accordance with the performance of the disk groups and the arrangement designated by the policy, and transmits the arranged disk groups to the program 20 for presenting the volume
5 candidate.

A detailed description is given of the flow for designating the condition for allocating the volume by the operating manager 10 and for actually allocating the volume by the storage management server 3 in the storage
10 management client 2 in order of processing 21, 311 to 318, and 22 to 23 shown in Fig. 1.

In processing 21 for designating the condition for allocating the condition, a policy list having a description of the condition for allocating the volume is
15 obtained from the policy database 35, displays the obtained policy list on a display screen of the storage management client 2, and presents it to the operating manager 10.

Fig. 7 shows a display example of a display screen for designating the condition for allocating the volume.

20 The display screen displays detailed contents of the policy, an index for selecting the policy, and an index for designating and inputting the memory capacity. Information on the policy is preset and is stored in the policy database 35.

25 The operating manager 10 selects the policy matching

the allocated volume from the policy list displayed on the display screen on the program 20 for presenting the volume candidate by use of the input device. In addition, the memory capacity of the allocated volume is designated. The 5 capacity of volume is designated to "20 GB" by, e.g., pulling down from the preset display. The program 20 for presenting the volume candidate transmits the selected and inputted policy and the designated volume capacity to the storage management server 3. The policy and volume 10 capacity may directly be inputted by the input device, not by a pull-down system.

Fig. 2 shows an example of a policy table stored in the policy database 35.

The policy comprises indexes of: a policy name 351; 15 peak performance 325; lowest-security performance 353; reliability 354; and operating time 355. The policy is previously formed by the operating manager 10 before starting the operation of volume, and is stored in the policy database 35. The policy name 351 is named by the 20 operating manager 10. The peak performance 352 is obtained by designating, by the operating manager 10, the performance to be secured at the time for operating the allocated volume. The lowest-security performance 353 is designated by the operating manager 10 when he desires the 25 security of the performance except for the operating time

of the allocated volume. The peak performance 352 and the lowest-security performance 353 are defined to be divided into read performances 3521 and 3531 and write performances 3522 and 3532, respectively.

5 The reliability 354 is performance of the storage for arranging the allocated volume. Upon defaulting the disk forming RAID, the restoring operation of data depends on the RAID level. Therefore, the reliability 354 is determined depending on the type of RAID. For example,
10 RAID 1 denotes "high", RAID 5 denotes "mid" (middle), and RAID 0 denotes "low".

The operating time 355 is determined by designating the time for operating the allocated volume on unit basis of time 3551, day 3552, week 3553, and month 3554. For
15 example, according to a policy 357 for weekday service, the read performance of the peak time is 50 MB/s, the write performance thereof is 10 MB/s, the read performance and the write performance except for the operating time are 10 MB/s, respectively, the reliability is "mid" (middle), and
20 the operating time is 9:00 to 17:00 on Monday to Friday.

A description returns to the functions of the program
31 for generating the volume candidate shown in Fig. 1.

In processing 311 for obtaining the disk group, the storage management server 3 forms a list of the disk group
25 which satisfies an input condition which is transmitted

from the program 20 for presenting the volume candidate of the storage management client 2. The list includes all the disk groups which are registered in the database 36 for the volume performance value. The database 36 for the volume 5 performance value stores information on the performance value of all the disk groups of the storage device 7 that can be accessed from the storage management server 3.

Fig. 3 shows an example of a table for performance-value information of the disk groups stored in the database 10 36 for the volume performance value.

The information of the disk groups includes a disk group name 361, reliability 362, a capacity 363, a theoretical performance value 364, and a forecasted performance value 365. The disk group name 361 indicates a 15 name of the disk group. The reliability 362 indicates a default ratio of the disk group. The capacity 363 comprises a maximum capacity 3631 and a used capacity 3832, indicating the maximum capacity and the currently used capacity of the disk group.

20 The program 33 for collecting the specification values calculates the theoretical performance value 364 based on the specification values 71 of the storage such as the performance of HDD, the number of drives, and the RAID level, and stores the calculated theoretical performance 25 value 364 in the database 36 for the volume performance

value. When the disk group of the RAID 0 comprises four drives having the performance of 50 MB/s, the theoretical performance value is 200 MB/s ($= 50 \times 4$). The theoretical performance value 364 comprises read performance 3641 and write performance 3642, and indicates the performance values upon reading time and writing time at which the disk group theoretically exhibits the maximum performance.

5 The forecasted performance value 365 comprises unit time 3651, read performance 3652, and write performance 10 3653, and indicates the forecasted performance value of the disk group per unit time. The forecasted performance value 365 is calculated from the information on the operation history and is obtained by calculating and forecasting the performance of the disk group upon allocating the volume 15 based on the performance history. Since the performance of the disk group is known by estimation, the performance of the disk group upon allocating the volume is not forecasted. Then, it is assumed that the performance of the disk group 20 upon allocating the volume is forecasted based on the information on the operation history of the storage. The information on the operation history is stored so that the actual estimated performance value of the disk group corresponds to the operating time.

Fig. 4 shows an example of a table for the information 25 on the operation history of the disk group stored in the

database 37 for the operation history.

Information of the disk group comprises a disk group name 371, and an actual estimated performance value 372. The disk group name 371 indicates the name of the disk group. The actual estimated performance 372 comprises time 3721, read performance 3722, and write performance 3723, and indicates time for estimating the performance of the disk group and an estimated value. The forecasted performance value 365 is obtained by segmenting the actual estimated performance value 372 per unit time and averaging it. For example, when the operation history is one month and the unit time is the day of week, the operation history is segmented per day of the week and is averaged per day of week. The forecasted performance value 365 and the unit time will be described in detail later. The reliability 362, capacity 363, and theoretical performance value 364 are obtained by collecting and processing the specification values of the storage device 7 on the program 33 for collecting the specification values and are stored in the database 36 for the volume performance value.

The list of the disk groups are obtained by comparing the disk groups under the following conditions (1) to (3) and by registering the disk groups which satisfy the entire conditions (1) to (3). The conditions (1) to (3) are as follows.

(1) Disk group having the reliability 362 more than the reliability designated as the input condition.

(2) Disk group in which the addition capacity of that designated as the input condition and of that used by the 5 existing volume in the disk group is not over the maximum capacity 363 thereof.

(3) Disk group having the theoretical performance 364 value which is not less than the performance designated as the input condition.

10 Fig. 5 shows an operation flowchart of the processing 312 for forecasting the disk-group performance value. The processing 312 for forecasting the disk-group performance value comprises processing 313 for extracting the information on the operation history of the disk groups 15 from the database 37 for the operation history and for segmenting data per unit time, and processing 314 for averaging the segmented data and for storing the average into the database 36 for the volume performance value as the forecasted performance value.

20 In the processing 313 for segmenting the information on the operation history per unit time, the unit time is obtained from the operating time zone of the designated policy (3131) and the information on the operation history is segmented per unit time (3132). If the operating time 25 zone is designated by hour unit, day unit, week unit, or

month unit, "hour", "day", "day of week", or "month" is the unit time, respectively. In an example of the policy 357 for "weekday service" shown in Fig. 2, the operation time zone is designated by the week unit and the unit time 5 is "day of week". A plurality of indexes are designated to the operating time zone and, in this case, the shortest time becomes the unit time. In an example of the "accounting month" 359 shown in Fig. 2, the unit time becomes "hour".

10 In processing 314 for averaging the segmented data and for storing the average, as the forecasted performance value, into the database 36 for the volume performance value, the average at the time is calculated based on the segmented information on the operation history and the 15 calculated average is obtained as the forecasted performance value (3141), and the forecasted performance value is stored in the database 36 for the volume performance value (3142).

20 A description is given of a method for obtaining the forecasted performance value by use of the policy 357 for "weekday service" shown in Fig. 2 and the volume performance value in the disk 1 shown in Fig. 3. In the designated policy, the operating time zone is designated by week unit. Consequently, in the processing 312 for 25 forecasting the disk-group performance value, the operation

history is segmented on each of Monday to Sunday, the performance value is averaged per day of week, and the forecasted performance value is stored in the database 36 for the volume performance value.

5 Fig. 6 shows an operation flowchart of the processing 315 for determining the allocated candidate volume. The processing 315 for determining the allocated candidate volume comprises processing 316 for obtaining a performance margin per unit time and processing 317 for deleting, from 10 the candidates, the disk group which does not satisfy the performance designated by the user.

The processing 316 for obtaining the performance margin per unit time is divided into processing 3161 for subtracting the forecasted performance value from the 15 theoretical performance value per unit time so as to obtain the performance margin and processing 3162 for subtracting the performance value designated by the policy from the performance margin per unit time. In the processing 3161 for subtracting the forecasted performance value from the theoretical performance value per unit time, the forecasted performance value is subtracted from the theoretical performance value 364 extracted from the database 36 for 20 the volume performance value so as to obtain the performance margin of the disk group per unit time. In the processing 3162 for subtracting the performance value 25

designated by the policy from the performance margin per unit time, the performance value designated by the policy is subtracted from the performance margin per unit time obtained in the processing 3161 for subtracting the 5 forecasted performance value from the theoretical performance value per unit time, and the performance value remaining in the disk group is obtained upon allocating the volume. The performance value designated by the policy includes two types of peak performance 342 and lowest-security performance 343. The performance value assured at 10 the operating time zone is designated to the peak performance 342 and the performance value secured as the lowest level is designated to the lowest-security performance 343 at the time except for the operating time 15 zone.

The processing 317 for deleting, from the candidates, the disk group which does not satisfy the performance designated by the user is divided into processing 3171 for determining whether or not all the subtracted values are 20 positive and processing 3172 for adding the disk group to the candidate list for allocating the volume.

In the processing 3171 for determining whether or not all the subtracted values are positive, it is determined whether or not positive are all the performance values 25 remaining in the disk group per unit time, which are

obtained in the processing 316 for obtaining the performance margin per unit time, or whether or not only one negative performance-value exists. If it is determined that all the performance values are positive, the disk
5 group is added to the candidate list for allocating the volume (3172). If it is determined that even one performance value is negative, the disk group is not added to the candidate list for allocating the volume and the processing 315 for determining the allocated candidate
10 volume of the remaining disk groups is repeated.

Next, a description is given of an example of calculating the performance margin in the processing 315 for determining the allocated candidate volume with reference to Fig. 9.

15 First, a description is given of a method for calculating the performance margin using the policy 357 for weekday service shown in Fig. 2 and the database 36 for the volume performance value shown in Fig. 3.

Through processing 3161 for subtracting the forecasted
20 performance value from the theoretical value per unit time to obtain the performance margin, a performance margin 801 per unit time is calculated. Further, a target performance value designated by the policy is subtracted from the performance margin 801, thereby obtaining the performance
25 value remaining in the disk group upon allocating the

volume. Since a disk 2 (8022) includes a negative value, it is understood that the disk 2 does not become the candidate for allocating the volume.

In the processing 318 for generating information 5 indicating the volume candidate, information is generated to display the volume candidate on the display screen of the storage management client 2. When a plurality of volume candidates exist in the candidate list for the volume, the volume candidates need to be arranged upon 10 displaying them on the display screen. The arrangement is performed based on the arrangement policy 356 in the policies (refer to Fig. 2) and the performance margin obtained in the processing 315 for determining the allocated candidate volume.

15 The arrangement policy 356 has two policies of concentration and distribution. The concentration means that the volumes are arranged to be concentrated to one disk group as much as possible so that the disk performance is effectively utilized. For example, the concentration is 20 used upon arranging the volumes having the low access-times (low frequency of access) to the single disk group in a lump. The distribution means that the volumes are arranged prior to an unused disk group. For example, the distribution is used upon distributing and arranging the 25 volumes having the high frequency of access to the

different disk groups. In the case of using the concentration as the arrangement policy 356, the volumes are arranged in order of the small sum of performance margins. In the case of using the distribution as the 5 arrangement policy 356, the volumes are arranged in order of the large sum of performance margins. The storage management server 3 transmits the candidate list for the volume to the storage management client 2.

In the processing 22 for presenting and selecting the 10 volume candidate, the volume-candidate list transmitted from the program 31 for generating the volume candidate is displayed on the display screen of the storage management client 2, and is presented to the operating manager 10.

Fig. 8 is a diagram showing an example of the screen 15 structure used in the processing 22 for presenting and selecting the volume candidate. The operating manager 10 instructs and selects, by use of the input device, one volume candidate from the candidate list for allocating the volume displayed on the display screen. If a request is 20 transmitted from the operating manager 10 in this case, the volume candidates can be sorted in order of the performance or of the reliability. In the processing 23 for allocating the volume, the selected volume candidate is actually allocated to the storage device 7.

25 The above-designated information on the candidate for

allocating the volume is transmitted to the storage management server 3 from the storage management client 2. On the program 34 for allocating the volume, the volume allocating processing is performed to the storage device 7.

5 As mentioned above, according to the embodiment, upon allocating the volume, the performance margin is calculated based on the performance value of the newly-allocated volume and on the performance value of the existing volume obtained from the information on the operation history, 10 thereby enabling the calculation of the volume as the allocated candidate so as to prevent its performance value which is over the maximum performance of the disk group. Consequently, the volume is efficiently and deliberately allocated, e.g., the volume operated on the weekday and the 15 volume operated on the holiday are allocated to the single storage device.

As mentioned above, the embodiment of the present invention is described. However, the present invention is not limited to this and is variously modified.

20 The table formats of the databases 35 to 37 are not limited to those described above. The index of the arrangement policy 356 is not necessary in the table of the policy database 35.

Further, in the table of the database 36 for the 25 volume performance value, the unit time in the forecasted

performance value 365 is not limited to month, week, day of week, and hour, and can be prescribed, e.g., every three days, every weekend, or every eight hours a day.

According to the present invention, advantageously, it
5 is possible to prevent the deterioration in performance which is caused by the excessive access performance of the allocated disk group when the volume is allocated and is then operated.